Bank of Canada Banque du Canada

Working Paper 2003-8 / Document de travail 2003-8

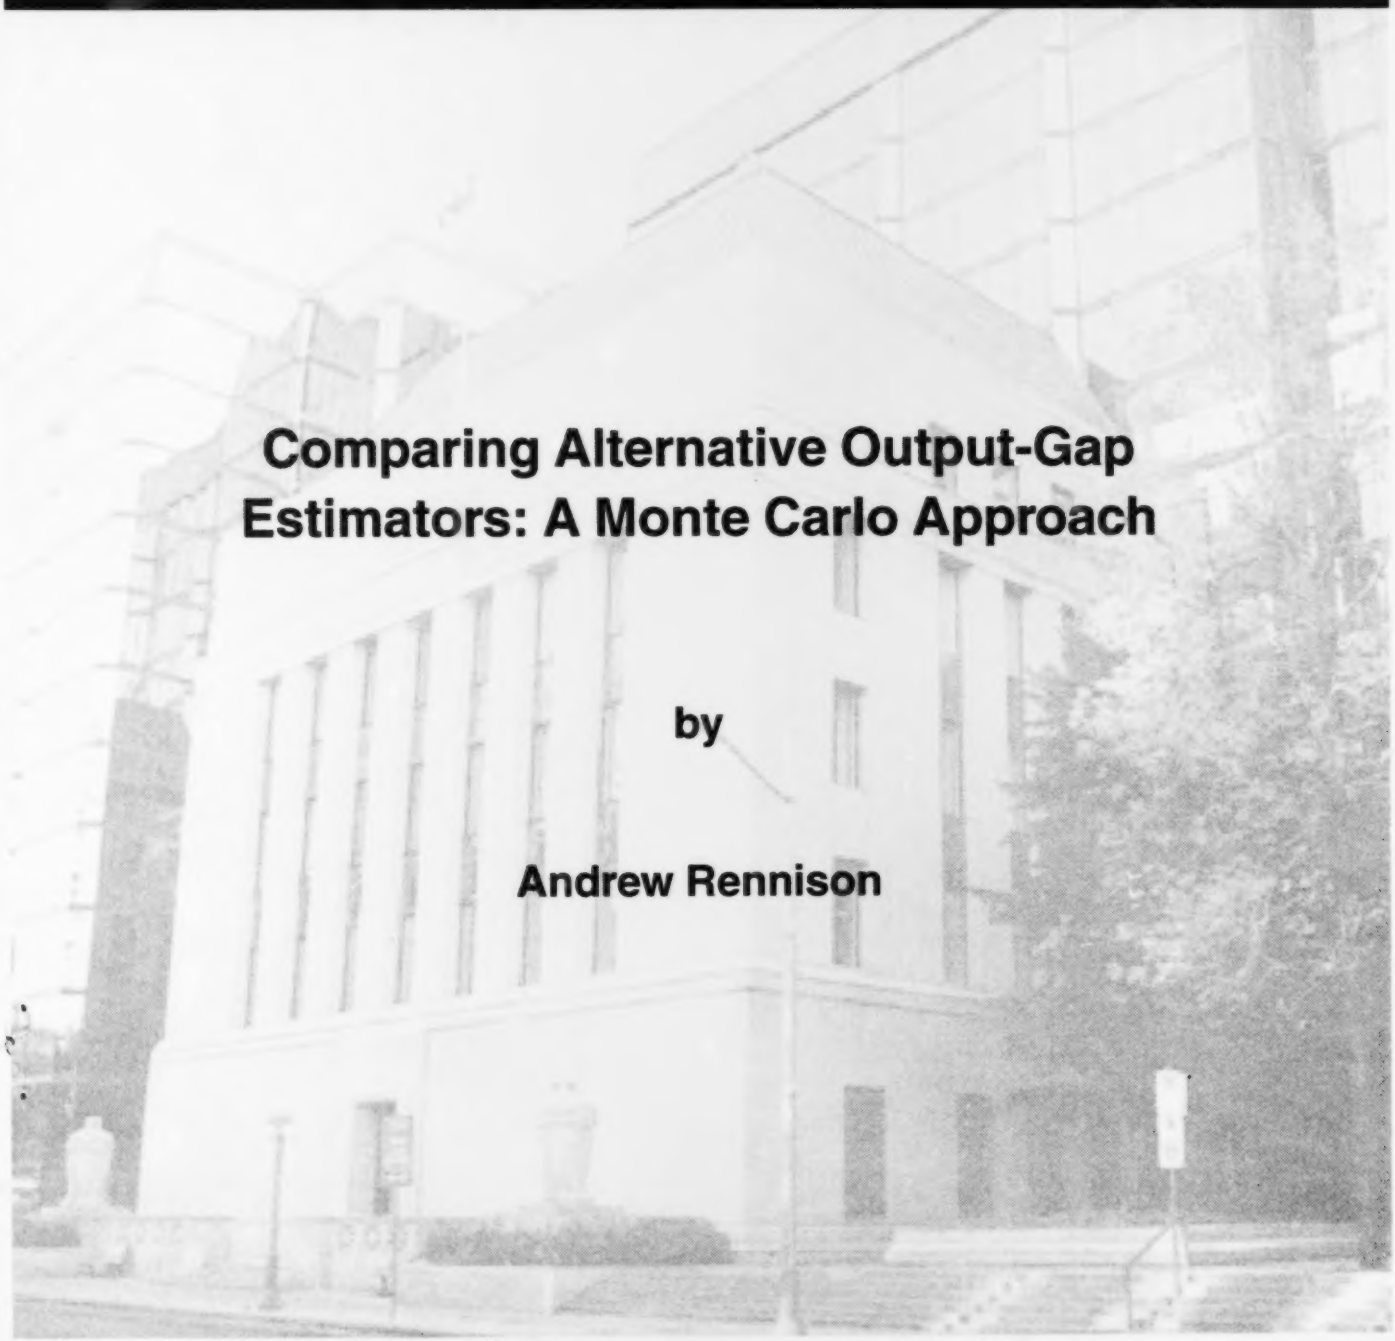

Comparing Alternative Output-Gap Estimators: A Monte Carlo Approach

by

Andrew Rennison

ISSN 1192-5434

Printed in Canada on recycled paper

103-05236/

Bank of Canada Working Paper 2003-8

March 2003

Comparing Alternative Output-Gap Estimators: A Monte Carlo Approach

by

Andrew Rennison

Research Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
arennison@bankofcanada.ca

The views expressed in this paper are those of the author.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements iv
Abstract/Résumé v

1. Introduction 1
2. The Data-Generating Process 2
3. The Methodologies 3
   - 3.1 The HP filter 4
   - 3.2 The multivariate HP filter 5
   - 3.3 The Blanchard-Quah SVAR 7
   - 3.4 The combined approach 8
4. The Experiment's Design 9
5. Results 11
   - 5.1 Correlations 13
   - 5.2 Error distributions 16
   - 5.3 RMSEs 18
6. Why the Combined Approach? 19
   - 6.1 Parameter uncertainty 20
   - 6.2 Model misspecification 20
7. Concluding Remarks 22

Bibliography 23

Acknowledgements

I would like to thank my collegues at the Bank of Canada, particularly Bob Amano, Don Coletti, René Lalonde, and Stephen Murchison for their helpful comments. Special thanks to Simon van Norden for sharing his expertise and enthusiasm.

Abstract

The author evaluates the ability of a variety of output-gap estimators to accurately measure the output gap in a model economy. A small estimated model of the Canadian economy is used to generate artificial data. Using output and inflation data generated by this model, the author uses each output-gap estimation methodology to construct an estimate of the true output gap. He then evaluates the methodologies by comparing their respective estimates of the output gap with the true gap. The estimators are evaluated on the basis of correlations between the actual and estimated output gap, as well as the root-mean-squared estimation error. The author also varies the properties of potential output and the output gap in the data-generating process to test the robustness of his results. His findings indicate that an estimator that combines the Hodrick-Prescott filter with a Blanchard-Quah structural vector autoregression (SVAR) yields an estimate that is accurate compared with competing methods at the end-of-sample. He also finds that the performance of the SVAR relative to that of other methodologies is quite robust to violations in the identifying assumptions of the SVAR.

JEL classification: C15, E32
Bank classification: Business fluctuations and cycles; Econometric and statistical methods; Potential output

Résumé

L'auteur évalue la capacité de différents estimateurs à mesurer l'écart de production dans le cadre d'une économie modélisée. Il utilise un petit modèle estimé de l'économie canadienne afin de générer des données artificielles sur la production et l'inflation. À partir de ces données, il s'attache à calculer le véritable écart de production en ayant recours à une batterie de méthodes d'estimation. Il analyse ensuite l'efficacité de ces dernières en comparant l'écart obtenu à l'aide de chacune d'elles au véritable écart. Les estimateurs sont évalués sur la base des corrélations entre l'écart effectif et l'écart estimé de même qu'à la lumière du critère de la racine carrée de l'erreur quadratique moyenne. Pour vérifier la robustesse des résultats, l'auteur modifie aussi les propriétés de la production potentielle et de l'écart de production dans le processus générateur des données. D'après ses conclusions, la combinaison du filtre de Hodrick-Prescott avec la méthode de Blanchard-Quah fondée sur l'emploi d'un vecteur autorégressif structurel donne une estimation plus précise que les autres méthodes en fin d'échantillon. L'auteur constate également que la méthode de Blanchard-Quah est supérieure aux autres même lorsque ses hypothèses d'identification ne sont pas respectées.

Classification JEL : C15, E32
Classification de la Banque : Cycles et fluctuations économiques; Méthodes économétriques et statistiques; Production potentielle

1. Introduction

It is generally accepted that the output gap—the difference between output and its potential or long-run sustainable level—is a key indicator of inflationary pressures and, as such, is an important variable for monetary policy. The construction of economic forecasts and the conduct of monetary policy are complicated, however, by the fact that potential output is unobservable and must therefore be estimated. Several competing methodologies exist for estimating the output gap, and there is a lack of consensus as to which is best. This paper evaluates some of the competing methodologies based on their ability to accurately measure the output gap in a model economy.

Because the output gap is unobservable, competing methodologies for estimating it are difficult to assess, and evaluation techniques have varied. Canova (1994), for example, uses the NBER definition of business cycle turning points as a metric for evaluating a battery of detrending methods. He finds that the Hodrick-Prescott (HP) (1997) filter does a good job, relative to other measures, of identifying turning points in U.S. real GDP. Recent work by Orphanides and van Norden (1999) for the United States, and by Cayen and van Norden (2002) for Canada, examines the sensitivity of several methodologies to the addition and revision of data at the end-of-sample, and finds that HP-filtered estimates of the Canadian output gap are subject to large real-time revisions compared with other methodologies. Alternatively, de Brouwer (1998) evaluates output-gap measures on the degree to which they help forecast inflation. Combining a simple forecasting equation with a variety of output-gap estimates, de Brouwer finds that the root-mean-squared error (RMSE) of inflation forecasts using a multivariate HP-filtered output gap is slightly smaller than that produced using various alternative gap measures. Most recently, however, Orphanides and van Norden (2001) have shown that, for a variety of methodologies, real-time estimates of the output gap provide little information in terms of out-of-sample inflation forecasting.

This paper takes a different approach, assessing some of the competing estimators of the output gap on the basis of their ability to accurately estimate the output gap of a model economy: Murchison's (2001) North American open economy macroeconometric integrated model (NAOMI). We focus on a subset of the available output-gap-estimation methodologies, specifically those used currently at the Bank of Canada, which generally are included in the family of HP-based or structural vector autoregression (SVAR) approaches. We also limit the set of information available to the various multivariate estimators; they use only data on output and inflation. Improvements in accuracy can be achieved by allowing a larger information set, and for this reason we hesitate to draw conclusions regarding the absolute accuracy of output gap estimators; we instead perform a relative assessment of the various estimators.

The evaluation takes the perspective of the monetary authority by focusing on the performance of the estimators at the end-of-sample. Correlations between the actual and estimated output gaps are used to gauge the ability of estimators to reproduce the dynamics of the output gap. We also consider the distributions of the estimation errors, focusing on the RMSEs.

The results of the evaluation suggest that the methodology that combines the multivariate HP filter with the SVAR is most robust across assumptions about (i) the relative volatility of the transitory and permanent components of GDP, (ii) the persistence of shocks to the growth rate of potential, (iii) the type of non-stationarity exhibited by potential output, and (iv) the degree of correlation between potential output and the output gap. The combined approach generally produces a more efficient (i.e., lower RMSE) estimate of the output gap at the end-of-sample. Perhaps most importantly, estimates produced by the combined approach are also generally the most highly correlated with the true output gap.

This paper is organized as follows. Section 2 briefly discusses the data-generating process, NAOMI. Section 3 describes and discusses the various estimation methods used, their key properties, and their underlying assumptions. Section 4 evaluates the various estimates on the basis of correlations and error distributions, and then further investigates the results for the combined approach. Section 5 concludes and briefly discusses some possible extensions.

2. The Data-Generating Process

NAOMI, the data-generating process (DGP) used for this investigation, is a small estimated model of the Canadian economy that consists of equations for the output gap, core CPI inflation, GDP inflation, real exchange rate, slope of the yield curve, long-term nominal interest rates, and, for the purposes of this paper, potential output.[1]

Although other models were considered, including the Bank's quarterly projection model (QPM), NAOMI was chosen for this analysis because of its simplicity and its ability to mimic well the historical dynamics of key variables in the Canadian economy. Also, importantly, the core of NAOMI is consistent with the central paradigm of models used in forecasting and projection at the Bank; monetary conditions affect the output gap via an IS curve, which in turn affects inflation via a Phillips curve. Potential output is defined in NAOMI as the level of output consistent with non-accelerating inflation.

1. Murchison (2001) provides a detailed description of NAOMI's properties.

Some changes to NAOMI were necessary for this study. First, an equation for potential output was added to its specification. In most of the experiments conducted in this paper, the growth rate of potential output is determined by an AR(1) process:

$$\Delta y_t^* = \mu + \delta \Delta y^*_{t-1} + \varepsilon_t, \tag{1}$$

where y_t^* is potential output and ε_t is an identically, independently distributed (i.i.d.) shock.[2] The value taken by the coefficient δ in the DGP is one dimension along which we can test the robustness of the output-gap estimators. Estimates of the degree of persistence in the growth of potential output will tend to vary with the method used to estimate potential output, and we therefore vary the value of this coefficient to test the robustness of our results. We also examine a case in which, rather than being difference stationary, potential output is stationary around a deterministic time trend:

$$y_t^* = \alpha + \beta t + \rho y^*_{t-1} + \varepsilon_t \,. \tag{2}$$

For each representation of potential output we ensure that the DGP is a plausible alternative representation of the true economy to the extent that it replicates two key observable features of the actual data: the volatility, as measured by the standard deviation, and persistence, as measured by the AR(1) coefficient, of first-differenced real GDP. Suppose, for example, that we conduct an experiment in which we increase the degree of persistence of potential output growth in the DGP. This change will, *ceteris paribus*, have a corresponding increase in the degree of persistence of overall output growth in the DGP. To offset this effect, we simply reduce the coefficient on the first difference of the output gap in the IS-equation such that the persistence in output growth is consistent with what we observe in the data. In a similar fashion, we ensure that the volatility of output growth in the artificial economy is equal to that in the actual data.[3]

3. The Methodologies

The methodologies considered in this paper are the HP filter and two multivariate techniques: the Blanchard-Quah (1989) SVAR approach and the multivariate extension of the HP filter (MVF). We also consider an estimator that is similar in spirit to the methodology used by the Bank's staff in that it weighs a portfolio of inputs to estimate the output gap.

---

2. This is equivalent to the form $y^*_t = \alpha + y^*_{t-1} + \nu_t$; $\nu_t = \delta \nu_{t-1} + \varepsilon_t$, where $\alpha = \mu \sum \delta^i$.
3. The estimated (1981 to 2001) historical standard deviation and AR(1) for first-differenced log GDP are 0.79 and 0.53, respectively.

Another popular method of estimating the output gap is the unobserved-component method, which includes the state-space model of, for example, Kuttner (1994). Unfortunately, owing to the typical instability of the maximum-likelihood estimates of the state-space model's parameters, we were unable to incorporate it into this study. In preliminary attempts to incorporate the model, the parameter estimates, in particular either the variance of potential output or output-gap shocks, tended towards zero in most samples. When dealing with one set of historical data, this problem can usually be overcome by trying various combinations of starting parameter values. In a Monte Carlo study such as this one, however, where estimates of the model are required for 10,000 samples of data, such an approach is virtually impossible.

3.1 The HP filter

The first estimator we examine is the simple, well-known, univariate HP filter. Although use of the HP filter and its variants remains widespread, the methodology has been subject to substantial criticism. In particular, its detractors point to the poor properties of HP-type filters at the end-of-sample, precisely where accuracy matters most for forecasting and policy decisions. St-Amant and van Norden (1997) examine the spectral properties of the HP filter and its multivariate extension, the MVF. They note the filter's inability to isolate business cycle frequencies in Canadian output data, particularly at the end-of-sample. Mise, Kim, and Newbold (2002) demonstrate that, while the HP filter is the optimal decomposition under certain conditions in mid-sample, at time-series endpoints it is suboptimal. Also, as stated in section 1, Cayen and van Norden (2002) use Canadian data to show that HP-filtered output gaps are extremely sensitive to data added at the end-of-sample. The HP filter's poor end-of-sample properties can be attributed to the fact that, while it is essentially a centred moving average in mid-sample, it becomes one-sided closer to the start and end of the sample. As such, the weight that the filter places on contemporaneous observations increases the closer the observation is to the end of the sample.

In the context of measuring potential output, the HP filter decomposes a time series for output into a cyclical component, the output gap, and a trend component, potential output. Application of the HP filter amounts to minimizing the variance of the cyclical component subject to a penalty for variance in the second difference of the trend component. Specifically, the HP filter solves for the value of y_t^*, which minimizes the following function:

$$\Theta_t = \sum_{t=0}^{T} (y_t - y_t^*)^2 + \lambda \sum_{t=1}^{T} (\Delta^2 y_t^*)^2 , \qquad (3)$$

where y_t is the raw series, y_t^* is the trend estimate, and λ is a smoothing prior.

Although the HP filter is generally viewed as an atheoretical method of detrending the data, its calibration can be given an economic interpretation. Specifically, the setting for λ, which controls the smoothness of the estimate of potential output, can be interpreted as a prior on the relative variance of supply and demand shocks. In the context of the current experiment, therefore, one would expect the HP filter to produce the most accurate estimates of the output gap when the DGP is consistent with this ratio. Alternatively, by varying the ratio of demand-to-supply shocks in our DGP, we can examine the costs of assuming the "wrong" value for λ.

3.2 The multivariate HP filter

The MVF combines the HP filter with at least one additional source of information (see Laxton and Tetlow 1992). Two versions of the MVF are examined in this study. The first (MVF1) adds the error from a pre-specified Phillips curve equation to the set of information used by the HP filter. Specifically, it chooses a profile for y_t^* that minimizes the following function:

$$\Theta_t = \sum_{t=0}^{T} (y_t - y_t^*)^2 + W_1 \sum_{t=4}^{T} (\varepsilon_t^{\pi})^2 + \lambda \sum_{t=1}^{T} (\Delta^2 y_t^*)^2 , \tag{4}$$

where ε_t^{π} is the error from a reduced-form Phillips curve relating the output gap to quarter-over-quarter core inflation. The user is required to specify the weight on the Phillips curve, W_1; the smoothing prior, λ; and the coefficients on the output gap in the Phillips curve, β_t.

The second version of the MVF (MVF2) is an extension of MVF1 that makes a correction of sorts for the end-of-sample problems associated with HP-based filters by stiffening the estimate of the trend series at the end-of-sample. Specifically, the minimization problem is extended as follows:

$$\Theta_t = \sum_{t=0}^{T} (y_t - y_t^*)^2 + W_1 \sum_{t=4}^{T} (\varepsilon_t^{\pi})^2 + W_2 \sum_{t=0}^{T} (y_t^* - y_t^{*P})^2 + W_3 (\Delta y_t^* - g)^2 + \lambda \sum_{t=1}^{T} (\Delta^2 y_t^*)^2 , \tag{5}$$

where y_t^{*P} is the estimate of the trend series from the previous quarter and g is an estimate of the steady-state growth rate of the trend series. So the filter is penalizing (i) the change in the estimated trend series as data are added at the end of sample, and (ii) penalizing deviations in the growth rate of the trend series from an estimate of its steady state.

Of course, the MVF estimate of the trend series is a function not only of the output and inflation data, but also of the user's choice of W_1, λ, and the set of β_ts. Unfortunately, within this

framework there is no formal way to choose values for these parameters in an optimal fashion.[4] Although the Phillips curve parameters can be estimated a priori, the choice of the weighting coefficient and the smoothing prior is more ad hoc. We use values for the weighting coefficients consistent with those used in the Bank's extended multivariate filter (see Butler 1996): λ equals 1600, the weighting coefficients W_1 and W_2 are set to one, and W_3 is set to 64 for the last 16 quarters of estimation and zero elsewhere.[5] g_{ss} is estimated simply as the mean growth rate of real output in the period of estimation.

An estimate of the Phillips curve parameters requires an initial estimate of the output gap. Following Conway and Hunt (1997), we use an HP-filtered output gap, with $\lambda = 1600$, as the initial estimate of the gap. Given that this experiment is being conducted from the standpoint of an economist who does not know with certainty the true structure of the economy, we impose the following general form for the Phillips curve:

$$\pi_t = \sum_{i=1}^{4} \alpha_i \pi_{t-i} + \sum_{i=0}^{4} \beta_i g_{t-i} + \varepsilon_t, \tag{6}$$

where π_t is inflation and g_t is the HP-filtered output gap. Inflation expectations are thus backward looking, and we impose that the coefficients β_i sum to one, consistent with the specification of the DGP. To eliminate the well-known end-of-sample problems of the HP filter, the first and last eight observations are eliminated from the estimation of the Phillips curve.

After the initial estimate of the Phillips curve parameters is obtained, the following iterative procedure is used to refine the output-gap estimate. First, the MVF estimates the output gap with the estimated Phillips curve coefficients. The Phillips curve coefficients are then re-estimated with this new output gap, and these coefficients are used to construct a new MVF estimate of the output gap. The procedure continues until the change in the output-gap estimate from one step to the next falls below a pre-specified convergence criterion.

---

4. Alternatively, this problem could be mapped into an unobserved-components model, in which the parameters are estimated by maximum likelihood.

5. See Butler (1996) for a discussion of the choice of these weights. Using estimates of the non-accelerating-inflation rate of unemployment (NAIRU) and the trend rate of capacity utilization, in addition to a Phillips curve, as conditioning information, de Brouwer (1998) sets the weights on each piece of conditioning information to be inversely proportional to the variance of the respective gap. Butler (1996), however, finds that such an approach does not produce estimates substantially different from those produced using a scheme of equal weights on each piece of conditioning information.

3.3 The Blanchard-Quah SVAR

The Blanchard-Quah (1989) SVAR methodology uses limited long-run restrictions to separate the temporary and permanent components of output. Fluctuations in output are attributed to two factors: those that have permanent effects on output, or supply shocks, and those that have temporary effects on output, or demand shocks. For this study, we use a bivariate SVAR with the first difference of log real GDP and the inflation rate, both of which are stationary in NAOMI.[6]

The long-run restriction used to identify the structural disturbances is that demand shocks have no long-run impact on the level of output.[7] We therefore impose the following long-run response matrix:

$$A(1) = \begin{bmatrix} y \\ \sum \pi \end{bmatrix} \begin{bmatrix} a_{11} & 0 \\ a_{21} & a_{22} \end{bmatrix}. \tag{7}$$

In other words, only one of the two shocks affects output in the long run, whereas both shocks can affect the *level* of prices in the long run (but not the inflation rate). This single restriction is sufficient, given the assumptions of orthogonality of the demand and supply shocks, to identify the structural shocks. The output gap is then computed as the cumulative response of the level of output to all past transitory shocks. The estimation starts with eight lags and tests down; first, a likelihood-ratio test is used to choose the lag length, after which the remaining coefficients whose *t*-statistics fall below one are discarded.[8]

It is important to note the consequences of the assumptions embodied in the SVAR methodology. First, the SVAR imposes an identifying restriction that demand and supply shocks are uncorrelated. Furthermore, as a consequence of this identifying assumption, the output gap and potential output are also uncorrelated, as the estimated output gap is simply the cumulation of all past demand shocks. There is therefore no channel whereby potential output is allowed to affect the output gap; the level of output adjusts one-for-one to a shock to the level of potential output.

---

6. Clearly, one would expect that additional information, such as a measure of monetary policy stance, for example, would aid the SVAR in identifying the structural shocks. It would be unsuitable, however, to compare the accuracy of a three-variable SVAR with a method such as the MVF, which uses data only on output and inflation.
7. Indeed, this long-run restriction holds in the DGP. Cooley and Dwyer (1998) show that a violation of the assumptions about the non-stationarity of the data can have a significant impact on how well an SVAR's dynamics mimic those in the true data. Future work will examine the implications of the trend growth rate of potential output being subject to structural breaks.
8. This approach to lag selection was proposed by Lutkepohl and Poskitt (1996).

This identifying assumption is at odds with QPM and NAOMI, both of which allow a channel whereby productivity shocks result in business cycle fluctuation. We gauge the consequences of incorrectly maintaining this assumption by comparing the cases in which there is no correlation between the transitory and permanent components of output with three cases in which correlation is allowed.

Cooley and Dwyer (1998) stress that the SVAR's "auxiliary" assumption of difference stationarity in output is not as innocuous as it appears. They point to the difficulty in distinguishing trend dependence from a unit root in post-war data to motivate an experiment in which output in the DGP, rather than being difference stationary, is instead driven by a near unit root, while the assumption of difference stationary output is still maintained in the SVAR. In this case, the identifying assumption that demand shocks have no long-run effect on output is no longer useful in separating demand and supply shocks. Cooley and Dwyer's results show that the dynamics implied by the SVAR dramatically distort those of the underlying data. Correspondingly, we simulate a version of NAOMI in which potential output is driven by persistent deviations around a time trend, to examine how much, if at all, the performance of the SVAR deteriorates relative to the other measures when this assumption is violated.

3.4 The combined approach

The Bank of Canada's methodology for estimating the output gap, the extended multivariate filter (EMVF), uses a portfolio of information to estimate the output gap. Specifically, the EMVF relies on several estimated reduced-form relationships in conjunction with an HP filter to estimate the gap. This approach is motivated by the fact that these reduced-form equations are prone to break down in the presence of structural change; by placing a weight on other relationships as well as the actual data, we are able to reduce the size of possible errors. An evaluation of the EMVF would require DGPs for variables beyond the scope of even QPM, and, as a result, we focus on a much simpler version, which captures the motivation behind the EMVF.[9]

This combined approach adds the SVAR estimate of the output gap, as conditioning information, to the MVF1 estimate. The MVF1 minimization problem is extended as follows:

9. For example, the EMVF estimate of the trend unemployment rate uses a structural VAR estimate of the NAIRU as conditioning information.

$$\Theta_t = \sum_{t=0}^{T} (y_t - y_t^*)^2 + W_1 \sum_{t=4}^{T} (\varepsilon_t^{\pi})^2 + W_2 \sum_{t=0}^{T} (y_t^* - y_t^{svar})^2 + \lambda \sum_{t=1}^{T} (\Delta^2 y_t^*)^2 , \qquad (8)$$

where y_t^{svar} is the estimate of potential output from the SVAR. The combined approach thus chooses the trend estimate, y_t^*, that minimizes (i) the difference between output and its trend series, (ii) the error in the Phillips curve, and (iii) the difference between the trend series and the SVAR estimate of potential, all subject to the smoothing prior. Consistent with the EMVF, the weights on the conditioning information, W_1 and W_2, are set to one, while the smoothing prior, λ, is set to 1600. An iterative procedure, identical to that used for the MVF1, is used to estimate the coefficients of the Phillips curve.

4. The Experiment's Design

Because of the lack of consensus on the underlying characteristics of potential output and the output gap, we consider several possibilities for the DGP. This is particularly important if we are to attempt to rank the various estimation methodologies, as the assumptions underlying certain estimators may bias our results for or against that technique. For example, identification of the SVAR requires an assumption about the type of non-stationarity in output data. On the other hand, it has been shown (King and Rebello 1993; Ehlgen 1998) that one condition under which the HP filter is optimal, in a mean-squared-error sense, is the smoothing parameter, λ, being equal to the ratio of the variances of innovations in the cyclical and trend components in the DGP. Changes in these characteristics in the DGP should therefore result in changes in the performance of the respective methodologies

Table 1 displays the characteristics of the various experiments. Column 2 shows the relative volatility, in terms of standard deviations, of the first differences of the output gap and potential output. For Case 1, and for most other cases, we set this ratio to 2, which is roughly consistent with the ratio typically yielded by HP-filtered estimates, with λ = 1600. Case 3 provides an example where this ratio is reversed, and potential output volatility dominates that of the gap.

Table 1: The DGP

Case	$sd(\Delta y^g)/sd(\Delta y^p)$	δ	ρ	ϕ [a]	$corr(y_t^g, \varepsilon_t^p)$
1	2	0.64	-	1.35	0
2	2	0.98	-	1.23	0
3	0.5	0.64	-	1.03	0
4	2	0.64	-	1.35	-0.18
5	2	0.98	-	1.25	-0.20
6	0.6	0.45	-	1.11	-0.42
7	2	-	0.95	1.55	0

a. Although coefficients on the lagged output gap are greater than one, other variables in the IS curve ensure that the output gap is, in fact, stationary in simulations of NAOMI.

The next two columns in Table 1 show the values taken by parameters controlling the persistence of potential output innovations in the difference stationary and trend stationary cases. Column 3 displays the coefficient δ, which controls the persistence of the first difference of potential output innovations when potential is defined by equation 1; for Cases 1 and 2 it takes the value of 0.64 and for Case 3 we increase it to 0.98.[10] Column 4 shows the value taken by the coefficient ρ in the case where potential output is determined by equation (2), violating the difference stationarity assumption of the SVAR. As stated earlier, to ensure that the properties of the DGP are as consistent as possible with the characteristics of historical data, we impose for each case that the volatility and persistence in the first difference of real GDP are equal to their historical values. Changing the persistence of potential output thus requires that the persistence of the output gap also be changed; this is done via the coefficient on the lagged output gap in the IS curve, the values of which are displayed in column 5.

The SVAR also makes the identifying assumption that innovations in the output gap and potential output are uncorrelated. For the first three experiments, we ensure this property in the DGP by removing potential output from NAOMI's output-gap equation, re-estimating the model,[11] and setting to zero the covariance between (permanent) potential output shocks and other (transitory)

10. 0.98 approximates the degree of persistence found in the Bank staff's EMVF estimate of potential, while 0.64 is the estimate obtained for a three-variable SVAR of the Canadian economy estimated from 1982 to 2001.

11. For the purposes of this paper, the change in parameter estimates is not large enough to merit discussion.

shocks. On the other hand, for Cases 4 through 6, we allow correlation between the permanent and transitory components, thereby violating the identifying assumption of the SVAR.[12] Case 6 is the most extreme: the only shocks are those to potential output, with all business cycles therefore caused by permanent shocks.[13] The final column in Table 1 shows the contemporaneous correlation between potential output shocks and the output gap.

For each experiment, 10,000 replications of the artificial-economy NAOMI are generated according to the estimated variance-covariance matrix of the model's residuals. For each of the resulting 10,000 samples of artificial data, an output-gap series is estimated by each methodology. Summary statistics for the resulting estimates' errors are then calculated within sample and averaged across samples. The estimation period is 160 quarters, or 40 years, which is similar in size to the typical sample of historical data. Since we are not overly concerned with results at the start of the estimation period, the full-sample summary statistics reported here are calculated excluding the first 28 quarters, which ensures the elimination of the start-of-sample problems associated with the various methodologies.[14]

5. Results

As an initial illustration, Figure 1 compares output-gap estimates from the MVF1, the SVAR, and the combined approach to the true output gap in one sample generated by the Case 1 parameterization of the DGP. In this sample, the three estimates follow the actual gap reasonably closely: the RMSEs for the MVF1, SVAR, and combined approach are 1.66, and 1.46, and 1.49, respectively, lower than the standard deviation of the actual output gap of 2.58. The correlations between the actual and estimated gap are 0.78, 0.86, and 0.83.

12. See Murchison (2001) for a discussion of supply shocks in NAOMI.
13. For Case 6, it is necessary to lower the coefficient on potential output growth to match the historical autocorrelation of output growth. Also, the ratio of demand-to-supply innovations of 0.6 is the lowest ratio possible, given the correlation between the output gap and potential in this case.
14. The start-of-sample (and/or end-of-sample) problems of the HP filter are well known and were described briefly earlier. The SVAR output gap is the cumulative output response to all past transitory shocks; at the start of the estimation period, when past shocks are unobservable, the cumulative impact of these shocks is of course also not estimable. When calculating summary statistics, we therefore cut off a portion of observations at the start of the estimation period, so that shocks that occurred before it will have for the most part died out in terms of their effect on output.

Figure 1: Output-Gap Estimates - Case 1

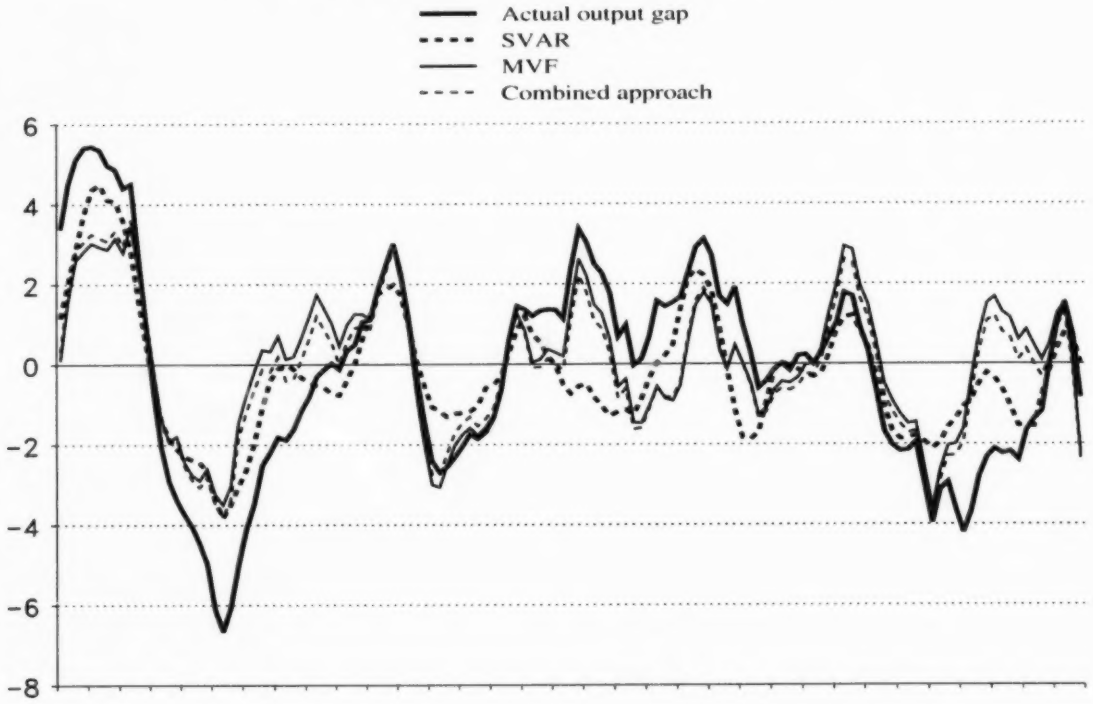

Alternatively, Figure 2 shows one sample from Case 3 in which all three estimators are at odds with the true data. The RMSEs for the MVF1, SVAR, and combined approach in this sample are 1.58, 2.62, and 1.66, respectively—larger than the standard deviation of the actual gap, 1.03. In other words, in an RMSE sense, an estimate of zero for all periods would dominate all three methods. It is important to recognize, however, that this result does not necessarily imply that the respective estimates of the gap will not be useful in deducing the state of the economy. Indeed, in this sample there exists a positive correlation between the actual gap and the MVF1 (0.33), the SVAR (0.74), and the combined approach (0.68). Perhaps most interesting is the contrast in results between the RMSE and correlation criteria; although the SVAR estimate has the highest RMSE, it is also the most highly correlated with the true output gap. It seems, therefore, at least in this particular sample, that the SVAR is able to match the sign and dynamics of the true output gap while failing to capture its correct magnitude. This example illustrates the important role multiple criteria play in an assessment of the performance of output-gap estimators.

Figure 2: Output-Gap Estimates - Case 3

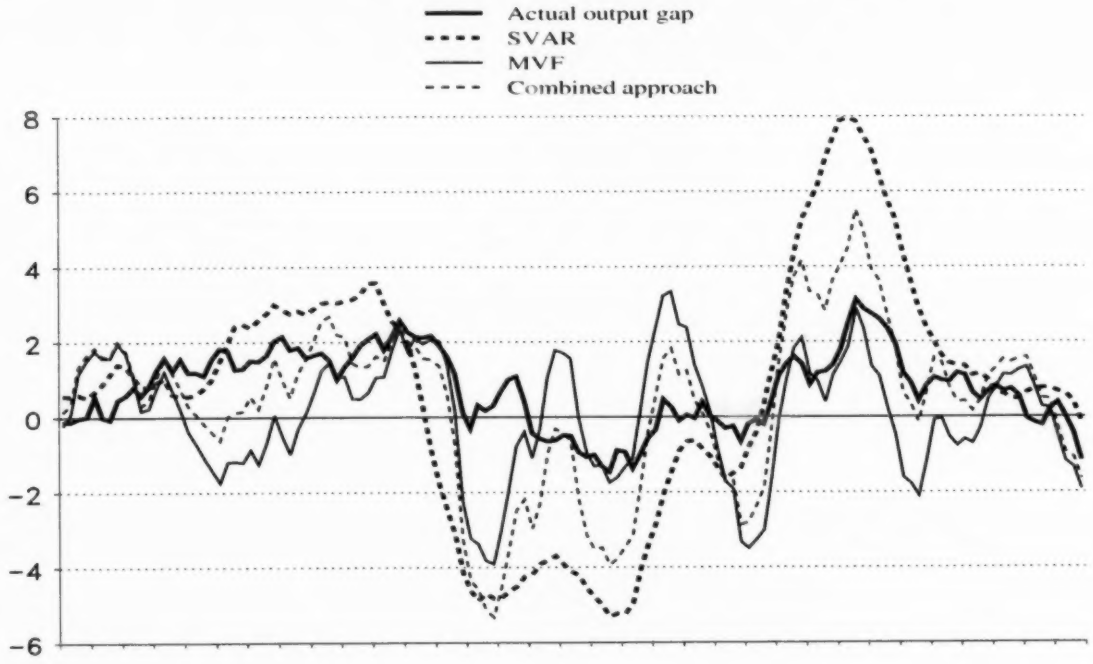

5.1 Correlations

It is easy to understand why, in some instances, it is more desirable for an output-gap estimator to have a high correlation between the estimated and actual output gap than a low RMSE. First, the correlation is more informative when evaluating an estimator's ability to correctly assess the state of the economy in terms of the sign of the output gap. Also, the correlation indicates the estimator's usefulness in forecasting inflation. For example, consider a hypothetical case in which an output-gap estimate is highly correlated with the true output gap but has a high RMSE (as with the SVAR estimate in Figure 2). Assuming the correct specification, the estimated Phillips curve parameters would simply be scaled up or down to compensate for the magnitude differences between the actual and estimated output gap. The accuracy of the inflation forecasts produced by such an output-gap estimate would therefore not be reflected by the inability to assess the magnitude of the gap, but rather by the degree to which the estimated gap is correlated with the true one. We do, however, examine the accuracy, in terms of RMSE, of the various estimators in section 4.2 under the assumption that, *ceteris paribus*, we would prefer an estimator that accurately assesses the magnitude of the output gap.

Table 2 shows the contemporaneous correlation coefficient between the estimated and actual gap for each of the estimators under the various DGPs, in full-sample and at the end-of-sample.[15] The results for a linear trend estimate of potential output are included for comparison. Several features of these results are noteworthy.

Table 2: Correlations

	Method	1	2	3	4	5	6	7
Full-sample	Linear	0.68	0.48	0.23	0.70	0.52	0.24	0.90
	HP	0.66	0.67	0.24	0.68	0.66	-0.22	0.80
	MVF1	0.66	0.67	0.24	0.69	0.67	-0.21	0.80
	MVF2	0.68	0.67	0.24	0.71	0.67	-0.14	0.82
	SVAR	0.58	0.58	0.27	0.54	0.43	0.05	0.65
	Combined	0.71	0.74	0.29	0.72	0.66	-0.15	0.83
End-of-sample	Linear	0.44	0.21	0.13	0.41	0.17	0.01	0.74
	HP	0.33	0.31	0.11	0.23	0.20	-0.89	0.43
	MVF1	0.33	0.32	0.12	0.24	0.22	-0.89	0.44
	MVF2	0.49	0.38	0.16	0.40	0.26	-0.42	0.65
	SVAR	0.43	0.40	0.20	0.38	0.29	0.02	0.48
	Combined	0.45	0.47	0.18	0.36	0.29	-0.51	0.53

We first examine the results for Cases 1 and 2, where the first difference of the output gap is roughly twice as volatile as the trend-growth component, potential output is difference stationary, and the orthogonality assumption of the SVAR holds in the DGP. The difference between the two cases is that potential output shocks are more persistent (and the output gap is thus less persistent) in Case 2 than in Case 1. The first obvious result is that changing the source of output persistence (i.e., moving from Case 1 to Case 2) does not greatly alter the correlation results. In both cases, the HP-based filters break down relative to the SVAR at the end of the sample. In mid-sample, the HP and multivariate HP (MVF1 and MVF2) filters produce output gaps that are more highly correlated with the true output gap than is the SVAR. On the contrary, at the end of the sample, the

15. In an initial investigation, the possibility of phase shift in the correlation structure was examined; under the Case 1 parameterization, the correlation structure for each estimator was symmetric, with a peak in contemporaneous observations. In this paper, the focus is therefore on the contemporaneous correlation.

relative performance of the SVAR improves. In both cases, the multivariate HP filter with the end-of-sample smoothing constraint (MVF2) outperforms the HP and MVF1. This result is not at all surprising when we consider that the end-of-sample constraints in the MVF2 would lead to a less volatile estimate of potential output. Overall, the best end-of-sample performers are the MVF2 for Case 1 and the combined approach for Case 2.

The results for Case 3, which is the same as Case 1 except that potential output is now more volatile than the output gap, reveal a significant deterioration in the performance of all approaches. Not surprisingly, the DGP's deviation from the two-to-one ratio of transitory-to-permanent fluctuations (the ratio typically yielded by the HP filter for Canadian real GDP) makes the performance of the HP and MVF filters deteriorate. This deterioration is such that the SVAR performs better at the end-of-sample, albeit by a small margin, than does the combined approach. Although the MVF2 also deteriorates at the end-of-sample, it again outperforms the other HP-based approaches. The SVAR performs the best at the end-of-sample.

For Cases 4, 5, and 6, the orthogonality assumption embedded in the SVAR is violated in NAOMI. Case 4 is simply Case 1 with correlation between the gap and potential, and Case 5 is the analogue of Case 2. Indeed, as expected, the performance of the SVAR deteriorates under these conditions. There is, however, a noticeable deterioration in the performance of the HP-based approaches, particularly at the end-of-sample. This phenomenon is not unexpected; King and Rebello (1993) discuss conditions under which the HP-filter is the optimal linear filter; they include orthogonal transitory and permanent innovations. As the DGP becomes less consistent with this condition, the performance of the HP filter worsens. It appears, then, that the consequences of correlation between demand and supply are similar for the SVAR and HP-based approaches. The performances of both have deteriorated to the point that, for Case 4, the linear trend produces the highest correlation at the end-of-sample. In Cases 4 and 5, the performance of the combined approach again stands out; in both full-sample and at the end-of-sample, the correlations for the combined approach are close to or higher than those for the SVAR and HP-based approaches.

The correlation results for Case 6, the pure potential output-shock case, are disconcerting. The SVAR output gap is uncorrelated with the true gap, whereas the HP-based estimates are negatively correlated with the true output gap (extremely so at the end-of-sample).[16] It would

16. To understand this result, first consider the dynamics of a supply shock in NAOMI: in a positive potential output shock, the output gap falls initially, while the level of output increases. The HP filter will interpret this increase in output partly as a positive demand shock and thus an increase in the output gap. This feature will be smaller in mid-sample, where the HP filter interprets the future observations of permanently higher output as evidence of a supply shock. To a lesser extent, the essentially zero correlation between the SVAR and the true output gap indicates that the SVAR is also attributing a sizable portion of the shocks to output as temporary rather than permanent.

appear that when fluctuations in the economy are driven solely by shocks to potential output, each of these output-gap estimation methods is extremely unreliable at identifying the state of the economy.

For Case 7, in which potential is driven by transitory deviations around a linear time trend, the results can be contrasted with those from Cases 1 and 2, in that transitory shocks dominate permanent ones while supply and demand are uncorrelated. Indeed, in terms of the relative performance of the various estimators, the results for Case 7 are very similar to those for Cases 1 and 2; the combined approach performs well in both mid-sample and full-sample, whereas the HP-based approaches break down at the end-of-sample. At the end-of-sample, the SVAR outperforms both the MVF1 and the HP filters. It is also not surprising that, in an economy with no permanent shocks to the level of potential output, the MVF2 performs well relative to other approaches. The same reasoning can be used to explain the superior performance of the linear trend. Overall, therefore, the results for Case 7 indicate that the costs of incorrectly maintaining the SVAR's assumption of difference stationarity in output are not large in terms of its ability to measure the output gap.

One striking aspect of the correlations is the generally strong performance of the combined approach at the end-of-sample. With the exception of Case 6, the combined approach produces the highest, or close to the highest, correlation with the true output gap. These results indicate that in many cases the estimates from the SVAR and the MVF1 are complementary, in that it is optimal to place some weight on the estimate from each approach rather than rely solely on one individual estimator. We will explore these results further in section 6.

5.2 Error distributions

This section examines the relative ability of the estimators to assess the magnitude of the output gap; that is, how does the size of estimation error compare across estimators? We also ask: do the results in terms of relative estimation accuracy agree with the correlation results of the previous section?

In the case of non-normality in the estimation errors, important information can be attained by examining the shape of the error distribution. For example, for two equally efficient estimators, there may be a larger probability of committing large errors with one estimator than with another, and such a possibility can be examined through the error distributions. Figures 3 through 6 show the end-of-sample error distributions for the HP, SVAR, MVF2,[17] and combined approach under the Case 1 parameterization. Clearly, little information about the accuracy of the estimators is contained in the tails of the error distributions. Indeed, this result holds across all experiments (for

17. The MVF1 and HP error distributions are virtually identical, so only those for HP are shown. We return to the issue of the similarity between the HP and MVF1 later in the paper.

brevity, we show only those for Case 1). The remainder of this section will therefore focus solely on the RMSEs.

Figure 3: Case 1 End-of-Sample HP Error Distribution

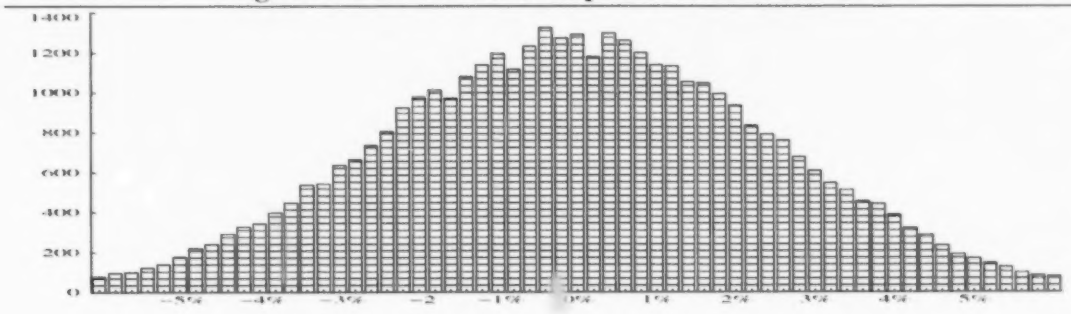

Figure 4: Case 1 End-of-Sample MVF2 Error Distribution

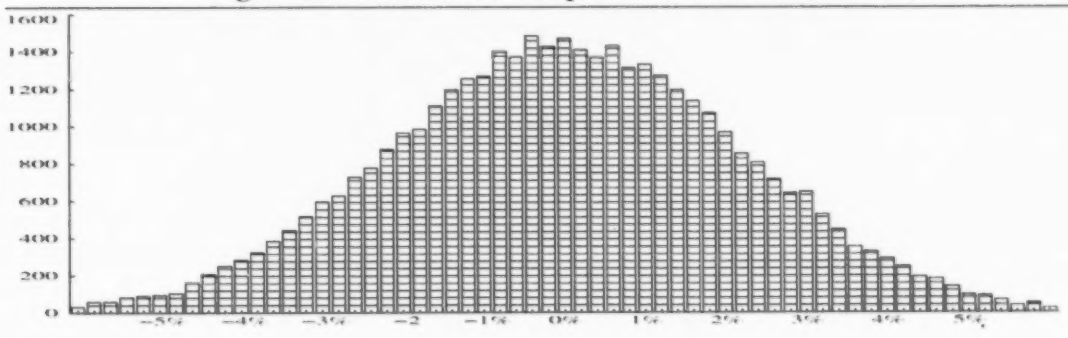

Figure 5: Case 1 End-of-Sample SVAR Error Distribution

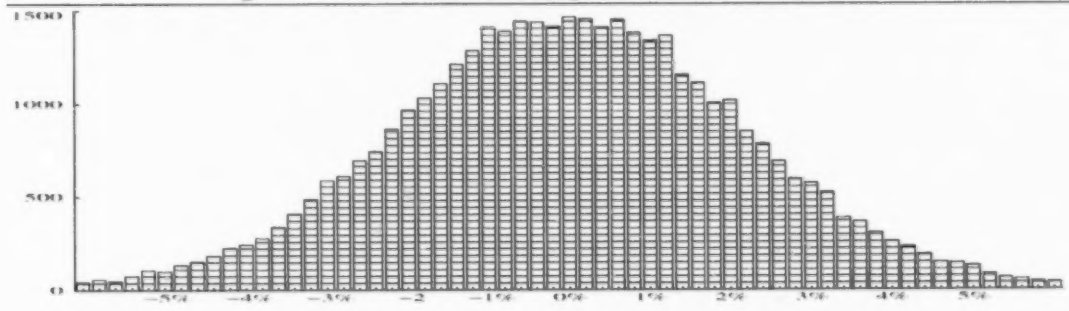

Figure 6: Case 1 End-of-Sample Combined-Approach Error Distribution

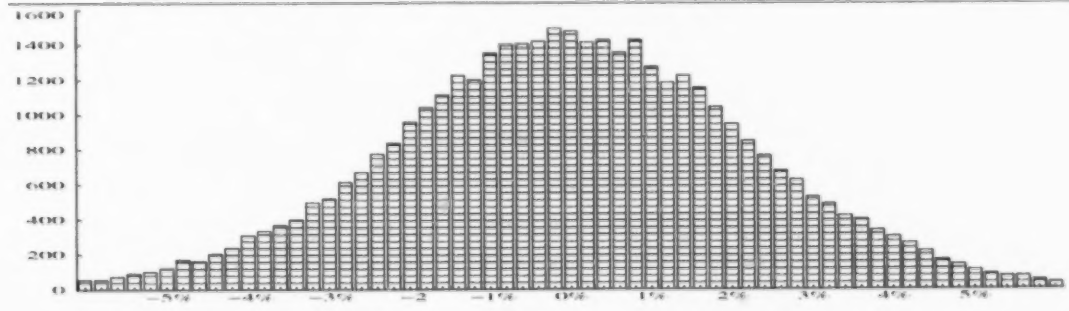

5.3 RMSEs

Overall, the RMSE statistics, shown in Table 3, conform broadly with the correlation results.[18] Again, the HP-based filters generally outperform the SVAR in mid-sample, and the reverse is true at the end-of-sample. Also, the combined approach stands out relative to the individual approaches; in both full-sample and at the end-of-sample, the combined output-gap estimate produces the lowest or close to the lowest RMSE. The outlier among the experiments is again Case 6, where fluctuations are driven by potential output shocks.

Table 3: RMSEs

	Method	1	2	3	4	5	6	7
Full-sample	Linear	2.47	4.78	4.63	2.69	4.92	7.36	1.10
	HP	1.93	1.97	1.71	2.16	2.22	3.57	1.50
	MVF1	1.92	1.96	1.71	2.15	2.20	3.56	1.50
	MVF2	1.88	1.95	1.78	2.09	2.17	3.61	1.44
	SVAR	2.07	2.21	1.93	2.39	2.59	5.31	1.85
	Combined	1.79	1.79	1.72	2.06	2.19	4.03	1.42
End-of-sample	Linear	3.46	8.25	6.44	3.96	8.86	10.89	1.39
	HP	2.52	2.56	1.75	2.99	3.00	4.44	2.33
	MVF1	2.50	2.53	1.74	2.97	2.97	4.44	2.24
	MVF2	2.19	2.68	2.41	2.68	3.35	5.46	1.69
	SVAR	2.22	2.46	2.00	2.59	2.82[a]	6.15	2.01
	Combined	2.25	2.26	1.69	2.72	2.83	4.67	2.04

a. The Diebold-Mariano statistic indicates that the Case 5 end-of-sample RMSEs for the SVAR and combined approach are not significantly different at the 10 per cent level. Otherwise, all other pairs of RMSEs are significantly different at 5 per cent.

18. Pairwise Diebold-Mariano (1995) statistics were calculated to assess the significance of the difference between end-of-sample RMSEs within each case. Except where noted, these differences are significant at the 5 per cent level.

One aspect in which the RMSE results are clearly distinct from the correlation results is the end-of-sample performance of the linear filter and the MVF2 filter. The shortcomings of the linear filter are now very obvious: the RMSE statistics for the linear filter at the end-of-sample are substantially higher than those for the other filters. To a lesser degree, the weakness of the MVF2 filter is illustrated by the results for Cases 2, 3, and 5, where, with the exception of the linear filter, it produces the highest RMSE. The RMSEs for the MVF2 filter reveal the potential costs associated with that filter's end-of-sample smoothing constraints. That is, in situations where trend movements in output are in fact persistent (as in Cases 2 and 5) or volatile (as in Case 3), the costs associated with penalizing changes in potential output growth can be rather large.

An additional aspect of these RMSE results is the remarkably similar performance of the HP and MVF1 (which is also readily apparent in section 5.1); the improvement in the MVF1's RMSE performance over that of the HP is very small for all seven cases in full-sample and at the end-of-sample. It is possible that this similarity is simply a result of the weight on the Phillips curve residual, W_1, in the MVF being too low, in which case the information from inflation would have very little influence on the estimate of potential output. There is, as stated earlier, no formal method for choosing a value for this coefficient; we therefore investigate this result further by increasing the value of W_1 tenfold and computing a new estimate of the MVF. In Case 1, the end-of-sample RMSE of the MVF1 output-gap estimate falls from 2.50 to 2.36, while its correlation with the true output gap rises from 0.33 to 0.40, indicating that the weight on the Phillips curve conditioning information used in the MVF may indeed be too low.

6. Why the Combined Approach?

The strong results for the combined approach are interesting; they indicate that estimates from the SVAR and the MVF1 are complementary (i.e., one estimator serves to offset the errors made by the other). Specifically, the MVF1 makes errors that can be reduced by putting some weight on the SVAR estimate of the output gap, and vice versa. An interesting question is posed by the end-of-sample results: why is it not optimal to rely entirely on the SVAR at the end-of-sample or, analogously, why is it optimal to put weight on the MVF1 at the end-of-sample where, as we have shown, its performance worsens? It must be the case that there is a sufficient amount of error in the SVAR estimate of the output gap that it is optimal to place some weight on an alternative, albeit flawed, estimator such as the MVF1. To understand this result it is instructive to think of a hypothetical case in which (i) the SVAR is correctly specified, and (ii) its parameters are precisely estimated. In this case, the SVAR's output-gap estimate would be exactly correct, and there would be no margin along which it could be improved. The further we move away from these conditions,

however, the less precise our SVAR estimate of the output gap becomes and, correspondingly, the more likely it becomes that the SVAR's estimate could be improved by relying in part on alternative estimators. In the following subsections, we therefore examine the consequences of parameter uncertainty and model misspecification in the SVAR.

6.1 Parameter uncertainty

The first possibility is parameter uncertainty associated with the SVAR. That is, perhaps our sample size is small enough that in some realizations of the model economy, poor estimates of the SVAR's parameters result in some weight being placed on the MVF1 at the end-of-sample. To investigate this possibility we rerun Case 1 using an estimation window of 500 periods, rather than the 160-period window used in the initial experiment.

Table 4 displays the results of this experiment and compares them with those for Case 1 at the end-of-sample. Notice that as we move from the 160-observation base case to the 500-observation case, the RMSE for the SVAR falls while the correlation rises. In the new case, the SVAR is the most favourable approach by both metrics. It appears that parameter uncertainty associated with the SVAR is at least partly responsible for the superior performance of the combined approach at the end-of-sample.

Table 4: End-of-Sample Performance, 500 Observations vs. 160 Observations

	160 observations		500 observations	
	Correlation	RMSE	Correlation	RMSE
MVF1	0.33	2.50	0.33	2.50
SVAR	0.43	2.22	0.51	1.98
Combined	0.46	2.25	0.48	2.18

6.2 Model misspecification

We investigate the possibility that misspecification in the SVAR may be the reason we find a role for the HP-filter at the end-of-sample. We do this in the context of Case 1, where the SVAR's assumptions of orthogonality and difference stationary output are both satisfied.

The easiest way to see how the SVAR is misspecified is to consider a version of NAOMI with shock terms only on the equations for potential output, the output gap (the IS curve), and inflation (the Phillips curve). Using a bivariate SVAR involves making the assumption that the economy is perturbed by only two "structural" shocks, one permanent and one transitory. Of course, even in this simplified version of NAOMI this is not the case; there are transitory shocks attached to both the IS curve and the Phillips curve and, furthermore, output and inflation respond in a different manner to each transitory shock. This example can be extended to include any number of variables; as long as each equation in the DGP is subject to random shocks, there will always be one more shock in the system than there are variables in the SVAR, and the SVAR will thus be misspecified.

We investigate this possibility by examining a case (Case 1b) in which the SVAR is not misspecified, in that the DGP contains only two shocks: (i) a shock to the Phillips curve, which has a transitory effect on output via the reaction function, and (ii) a shock to potential output, which has a long-run impact on the level of potential. Thus, relative to Case 1, we have removed the shocks to all equations except those for potential output and the Phillips curve.

Table 5 reports the results of this experiment relative to Case 1. Again, the SVAR is now the superior approach by both the RMSE and correlation criteria. Clearly, misspecification of the SVAR plays a role in determining the weight one should put on the SVAR at the end-of-sample.

Table 5: End-of-Sample Performance, Case 1 vs. 1b

	Case 1		Case 1b	
	Correlation	RMSE	Correlation	RMSE
MVF1	0.33	2.50	0.20	1.82
SVAR	0.43	2.22	0.65	1.17
Combined	0.46	2.25	0.54	1.30

7. Concluding Remarks

This paper has described a method for assessing the relative usefulness of a variety of output-gap measures using simulated data from an artificial economy. It has highlighted the potential benefits of using an approach that combines the HP-based and SVAR methodologies. This combined approach generally provides the most useful estimate of the gap in an RMSE sense and in terms of correlation between the actual and estimated output gap, indicating that the output-gap estimates from the SVAR and the HP-based filter are in many cases complementary. Our results appear quite robust to alternative realistic assumptions about the DGP. We have shown that the favourable results for the combined approach at the end-of-sample are due in part to misspecification and parameter uncertainty in the SVAR.

Two additional results have been reported: (i) relative to other estimation methodologies, the SVAR is surprisingly robust to violations in its identifying assumptions, and (ii) in terms of the absolute accuracy of an estimator at the end-of-sample, the costs associated with imposing an arbitrary smoothing restriction can be high.

There are several possible extensions to this study. Future work will examine the possibility of structural change, particularly in the growth rate of potential output. In other words, how do the various approaches perform when there is, say, a U.S.-style new-economy scenario, in which the trend growth rate of productivity increases (or may have increased)? It would also be interesting to attempt to determine a weighting scheme for the various estimates of the gap that is optimal in some sense.

Of course, as stated earlier, the results presented in this paper are no doubt partly a function of the DGP. Although several modifications to the DGP were considered in this study, future work could examine the robustness of our results to variation in the structure of the model economy.

Bibliography

Beveridge, S. and C.R. Nelson. 1981. "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle'." *Journal of Monetary Economics* 7: 151–74.

Blanchard, O.J. and D. Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review* 79(4): 655–73.

Butler, L. 1996. *A Semi-Structural Method to Estimate Potential Output: Combining Economic Theory with a Time-Series Filter.* Technical Report No. 77. Ottawa: Bank of Canada.

Canova, F. 1994. "Detrending and Turning Points." *European Economic Review* 38: 614–23.

———. 1998. "Detrending and Business Cycle Facts." *Journal of Monetary Economics* 41: 475–512.

Cayen, J-P. and S. van Norden. 2002. "La fiabilité des estimations de l'écart de production au Canada." Bank of Canada Working Paper No. 2002-10.

Clark, P.K. 1987. "The Cyclical Component of U.S. Economic Activity." *Quarterly Journal of Economics* 102(4): 797–814.

Conway, P. and B. Hunt. 1997. "Estimating Potential Output: A Semi-Structural Approach." Reserve Bank of New Zealand Discussion Paper No. G97/9.

Cooley, T.F. and M. Dwyer. 1998. "Business Cycle Analysis Without Much Theory: A Look at Structural VARs." *Journal of Econometrics* 83: 57–88.

de Brouer, G. 1998. "Estimating Output Gaps." Reserve Bank of Australia Research Discussion Paper No. 9809.

Diebold, F.X. and R.S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13(3): 253–63.

Ehlgen, J. 1998. "Distortionary Effects of the Optimal Hodrick-Prescott Filter." *Economics Letters* 61: 345–49.

Gerlach, S. and F. Smets. 1997. "Output Gaps and Inflation: Unobservable Components Estimates for the G-7 Countries." Bank for International Settlements. Basel. Photocopy.

Guay, A. and P. St-Amant. 1996. *Do Mechanical Filters Provide a Good Approximation of Business Cycles?* Technical Report No. 78. Ottawa: Bank of Canada.

Harvey, A.C. and A. Jaeger. 1993. "Detrending and Stylized Facts, and the Business Cycle." *Journal of Applied Econometrics* 8: 231–47.

Hodrick, R.J. and E.C. Prescott. 1997. "Postwar U.S. Business Cycles." *Journal of Money, Credit and Banking* 29: 1–16.

Kichian, M. 1999. "Measuring Potential Output Within a State-Space Framework." Bank of Canada Working Paper No. 99-9.

King, R.G. and S.T. Rebello. 1993. "Low-Frequency Filtering and Real Business Cycles." *Journal of Economic Dynamics and Control* 17: 207–31.

Kuttner, K.N. 1994. "Estimating Potential Output as a Latent Variable." *Journal of Business and Economic Statistics* 12: 361–68.

Laxton, D. and R. Tetlow. 1992. *A Simple Multivariate Filter for the Measurement of Potential Output.* Technical Report No. 59. Ottawa: Bank of Canada.

Lutkepohl, H. and D.S. Poskitt. 1996. "Testing for Causation Using Infinite Order Vector Autoregressive Processes." *Econometric Theory* 7: 61–87.

Mise, E., T.-H. Kim, and P. Newbold. 2002. "The Hodrick-Prescott Filter at Time Series End-Points." Manuscript.

Murchison, S. 2001. "NAOMI: A New Quarterly Forecasting Model Part II: A Guide to Canadian NAOMI." Department of Finance Working Paper No. 2001-25.

Orphanides, A. and S. van Norden. 1999. "The Reliability of Output Gap Estimates in Real Time." Board of Governors of the Federal Reserve System Finance and Economics Discussion Series No. 1999-38.

———. 2001. "The Reliability of Inflation Forecasts Based on Output Gaps in Real Time." *Computing in Economics and Finance* No. 247.

St-Amant, P. and S. van Norden. 1997. *Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada.* Technical Report No. 79. Ottawa: Bank of Canada.

Watson, M.W. 1986. "Univariate Detrending Methods with Stochastic Trends." *Journal of Monetary Economics* 18: 49–75.

Wecker, W. 1979. "Predicting the Turning Points of a Time Series." *Journal of Business* 52: 35–50.

Whittle, P. 1983. *Prediction and Regulation by Linear Least-Square Methods.* Minneapolis, MN: University of Minnesota Press.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. ***Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.***

2003

2003-7	Testing the Stability of the Canadian Phillips Curve Using Exact Methods	L. Khalaf and M. Kichian
2003-6	Valuation of Canadian- vs. U.S.-Listed Equity: Is There a Discount?	M.R. King and D. Segal
2003-5	Shift Contagion in Asset Markets	T. Gravelle, M. Kichian, and J. Morley
2003-4	Are Distorted Beliefs Too Good to be True?	M. Misina
2003-3	Modélisation et prévision du taux de change réel effectif américain	R. Lalonde and P. Sabourin
2003-2	Managing Operational Risk in Payment, Clearing, and Settlement Systems	K. McPhail
2003-1	Banking Crises and Contagion: Empirical Evidence	E. Santor

2002

2002-42	Salaire réel, chocs technologiques et fluctuations économiques	D. Tremblay
2002-41	Estimating Settlement Risk and the Potential for Contagion in Canada's Automated Clearing Settlement System	C.A. Northcott
2002-40	Inflation Changes, Yield Spreads, and Threshold Effects	G. Tkacz
2002-39	An Empirical Analysis of Dynamic Interrelationships Among Inflation, Inflation Uncertainty, Relative Price Dispersion, and Output Growth	F. Vitek
2002-38	Oil-Price Shocks and Retail Energy Prices in Canada	M. Chacra
2002-37	Alternative Public Spending Rules and Output Volatility	J.-P. Lam and W. Scarth
2002-36	Une approche éclectique d'estimation du PIB potentiel américain	M.-A. Gosselin and R. Lalonde
2002-35	The Impact of Common Currencies on Financial Markets: A Literature Review and Evidence from the Euro Area	L. Karlinger

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca